# Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data

**Jean Honorio,** **Luis Ortiz**
Department of Computer Science
Stony Brook University
Stony Brook, NY 11794
{jhonorio,leortiz}@cs.sunysb.edu

## Abstract

We formalize and study the problem of learning the structure and parameters of *graphical games* from strictly *behavioral* data. We cast the problem as a maximum likelihood estimation (MLE) based on a generative model defined by the *pure-strategy Nash equilibria (PSNE)* of the game. The formulation brings out the interplay between goodness-of-fit and model complexity: good models capture the equilibrium behavior represented in the data while controlling the *true* number of PSNE, including those potentially unobserved. We provide a generalization bound for MLE. We discuss several optimization algorithms including *convex loss minimization (CLM)*, sigmoidal approximations and exhaustive search. We formally prove that games in our hypothesis space have a small *true* number of PSNE, with high probability; thus, CLM is sound. We illustrate our approach, show and discuss promising results on synthetic data and the U.S. congressional voting records.

## 1 Introduction

*Graphical games* [Kearns et al., 2001] were one of the first and most influential graphical models for game theory. It has been about a decade since their introduction to the AI community. There has also been considerable progress on problems of *computing* classical equilibrium solution concepts such as Nash [Nash, 1951] and correlated equilibria [Aumann, 1974] in graphical games (see, e.g., Kearns et al. [2001], Vickrey and Koller [2002], Ortiz and Kearns [2002], Blum et al. [2006], Kakade et al. [2003], Papadimitriou and Roughgarden [2008], Jiang and Leyton-Brown [2011] and the references therein). Indeed, graphical games played a prominent role in establishing the computational complexity of computing Nash equilibria in general normal-form games (see, e.g., Daskalakis et al. [2009] and the references therein).

Relatively less attention has been paid to the problem of *learning* the structure of graphical games from data. Addressing this problem is essential to the development, potential use and success of game-theoretic models in practical applications.

Indeed, we are beginning to see an increase in the availability of data collected from processes that are the result of deliberate actions of agents in complex system. A lot of this data results from the interaction of a large number of individuals, being people, companies, governments, groups or engineered autonomous systems (e.g. autonomous trading agents), for which any form of global control is usually weak. The Internet is currently a major source of such data, and the smart grid, with its trumpeted ability to allow individual customers to install autonomous control devices and systems for electricity demand, will likely be another one in the near future.

We present a formal framework and design algorithms for learning the structure and parameters of graphical games [Kearns et al., 2001] in large populations of agents. We concentrate on learning from purely behavioral data. We expect that, in most cases, the parameters quantifying a utility function or best-response condition are unavailable and hard to determine in real-world settings. The availability of data resulting from the observation of an agent *public behavior* is arguably a weaker assumption than the availability of agent *utility* observations, which are often *private*.

Our technical contributions include a novel generative model of behavioral data in Section 2 for general games. We define identifiability and triviality of games. We provide conditions which ensures identifiability among non-trivial games. We then present the maximum likelihood estimation (MLE) problem for general (non-trivial identifiable) games. In Section 3, we show a generalization bound for the MLE problem and an upper bound of the VC-dimension of LIGs. In Section 4, we approximate the original problem by maximizing the number of observed equilibria in the data, suitable for a class of games with small *true* number of equilibria. We then present our convex loss minimization approach, and baseline methods such as sigmoidal approximation and exhaustive search for linear influence games (LIGs). In Section 5, we define absolute-indifference of players and show that convex loss minimization produces games in which all players are non-absolutely-indifferent. We provide a distribution-free bound which shows that LIGs have small *true* number of equilibria with high probability.

**Related Work.** Our work *complements* the recent line of work on learning graphical games [Vorobeychik et al., 2005, Ficici et al., 2008, Duong et al., 2009, Gao and Pfeffer, 2010, Ziebart et al., 2010, Waugh et al., 2011]. With the exception of Ziebart et al. [2010], Waugh et al. [2011], previous methods assume that the actions as well as corresponding payoffs (or noisy samples from the true payoff function) are observed in the data. Another notable exception is a recently proposed framework from the learning theory community to model *collective* behavior [Kearns and Wortman, 2008]. The approach taken there considers dynamics and is based on stochastic models. Our work differs from methods that assume that the game is known [Wright and Leyton-Brown, 2010]. The work of Vorobeychik et al. [2005], Gao and Pfeffer [2010], Wright and Leyton-Brown [2010], Ziebart et al. [2010] present experimental validation mostly for 2 players only, 7 players in Waugh et al. [2011] and up to 13 players in Duong et al. [2009].

In this paper, we assume that the joint-actions is the only observable information. To the best of our knowledge, we present the first techniques for learning the structure and parameters of large-population graphical games from joint-actions only. Furthermore, we present experimental validation in games of up to 100 players. Our convex loss minimization approach could potentially be applied to larger problems since it is polynomial-time.

**Background.** In classical game-theory (see, e.g. Fudenberg and Tirole [1991] for a textbook introduction), a *normal-form game* is defined by a set of *players* $V$ (e.g. we can let $V = \{1, \dots, n\}$ if there are $n$ players), and for each player $i$, a set of *actions*, or *pure-strategies* $A_i$, and a payoff function $u_i : \times_{j \in V} A_j \to \mathbb{R}$ mapping the joint-actions of all the players, given by the Cartesian product $\mathcal{A} \equiv \times_{j \in V} A_j$, to a real number. In non-cooperative game theory we assume players are greedy, rational and act independently, i.e. each player $i$ always want to maximize their own utility, subject to the actions selected by others, irrespective of how the optimal action chosen help or hurt others.

A core solution concept in non-cooperative game theory is that of an *Nash equilibrium*. A joint-action $\mathbf{x}^* \in \mathcal{A}$ is a *pure-strategy Nash equilibrium (PSNE)* of a non-cooperative game if, for each player $i$, $x_i^* \in \arg\max_{x_i \in A_i} u_i(x_i, \mathbf{x}_{-i}^*)$; that is, $\mathbf{x}^*$ constitutes a *mutual best-response*, no player $i$ has any incentive to unilaterally deviate from the prescribed action $x_i^*$, given the joint-action of the other players $\mathbf{x}_{-i}^* \in \times_{j \in V - \{i\}} A_j$ in the equilibrium.

In what follows, we denote a game by $\mathcal{G}$, and the PSNE set of $\mathcal{G}$ by

$\mathcal{NE}(\mathcal{G}) \equiv \{\mathbf{x}^* \mid \forall i \in V,\ x_i^* \in \arg\max_{x_i \in A_i} u_i(x_i, \mathbf{x}_{-i}^*)\}$.

A *(directed) graphical game* is a game-theoretic graphical model [Kearns et al., 2001]. It provides a succinct representation of normal-form games. In a graphical game, we have a (directed) graph $G = (V, E)$ in which each node in $V$ corresponds to a player in the game. The interpretation of the edges/arcs $E$ of $G$ is that the payoff function of player $i$ is only a function of the set of parents/neighbors $\mathcal{N}_i \equiv \{j \mid (i,j) \in E\}$ in $G$ (i.e. the set of players corresponding to nodes that point to the node corresponding to player $i$ in the graph). In the context of a graphical game, we refer to the $u_i$'s as the *local payoff functions/matrices*.

*Linear influence games (LIGs)* [Irfan and Ortiz, 2011] are a sub-class of graphical games. For LIGs, we assume that we are given a matrix of influence weights $\mathbf{W} \in \mathbb{R}^{n \times n}$, with $\mathbf{diag}(\mathbf{W}) = \mathbf{0}$, and a threshold vector $\mathbf{b} \in \mathbb{R}^n$. For each player $i$, we define the influence function $f_i(\mathbf{x}_{-i}) \equiv \sum_{j \in \mathcal{N}_i} w_{ij} x_j - b_i = \mathbf{w}_{i,-i}{}^{\mathrm{T}} \mathbf{x}_{-i} - b_i$ and the payoff function $u_i(\mathbf{x}) \equiv x_i f_i(\mathbf{x}_{-i})$. We further assume binary actions: $A_i \equiv \{-1, +1\}$ for all $i$. The *best response* $x_i^*$ of player $i$ to the joint-action $\mathbf{x}_{-i}$ of the other players is defined as

$$\begin{aligned}
&\mathbf{w}_{i,-i}{}^{\mathrm{T}} \mathbf{x}_{-i} > b_i \Rightarrow x_i^* = +1, \\
&\mathbf{w}_{i,-i}{}^{\mathrm{T}} \mathbf{x}_{-i} < b_i \Rightarrow x_i^* = -1 \text{ and} \\
&\mathbf{w}_{i,-i}{}^{\mathrm{T}} \mathbf{x}_{-i} = b_i \Rightarrow x_i^* \in \{-1, +1\},
\end{aligned}$$

or equivalently $x_i^*(\mathbf{w}_{i,-i}{}^{\mathrm{T}} \mathbf{x}_{-i} - b_i) \geq 0$. Hence, for any other player $j$, $w_{ij} \in \mathbb{R}$ can be thought as a *weight* parameter quantifying the "influence factor" that $j$ has on $i$, and $b_i \in \mathbb{R}$ as a *threshold* parameter to the level of "tolerance" that player $i$ has for playing $-1$.

Figure 1 provides a preview illustration of the application of our approach to congressional voting.

## 2 Problem Formulation

Our goal is to learn the structure and parameters of a graphical game from observed joint-actions. Our problem is unsupervised, i.e. we do not know a priori which

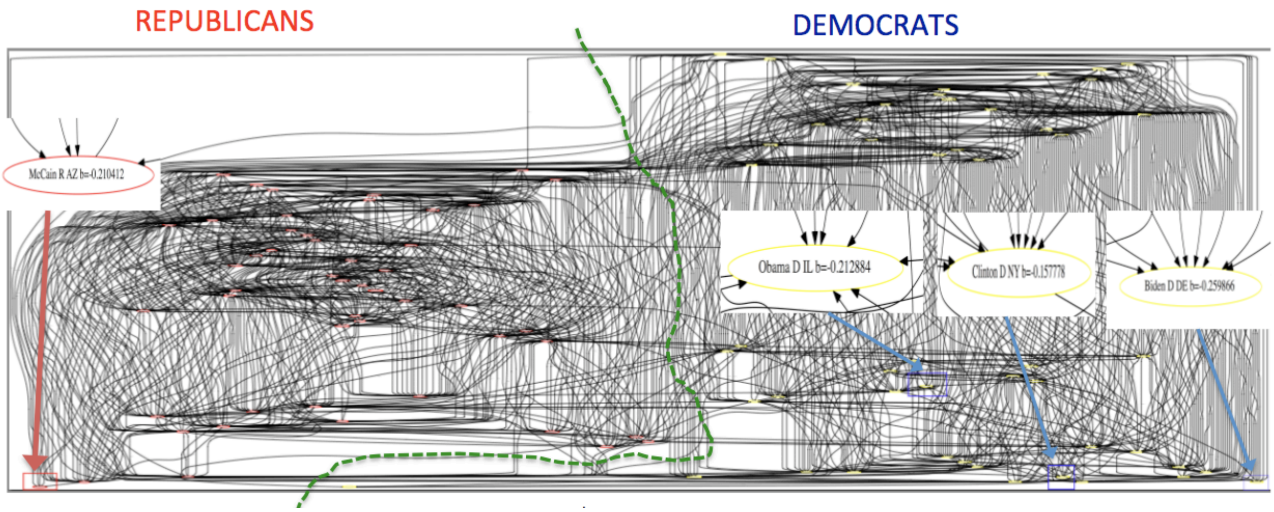


Figure 1: **110th US Congress's Linear Influence Game (January 3, 2007-09):** We provide an illustration of the application of our approach to real congressional voting data. Irfan and Ortiz [2011] use such LIGs to address a variety of computational problems, including the identification of *most influential* senators. We show the graph connectivity of a LIG learnt by independent $\ell_1$-regularized logistic regression (see Sect. 4.4). We highlight some characteristics of the graph, consistent with anecdotal evidence. First, senators are more likely to be influenced by members of the same party than by members of the opposite party (the dashed green line denotes the separation between the parties). Republicans were "more strongly united" (tighter connectivity) than Democrats at the time. Second, the current US Vice President Biden (Dem./Delaware) and McCain (Rep./Arizona) are displayed at the "extreme of each party" (Biden at the bottom-right corner, McCain at the bottom-left) eliciting their opposite ideologies. Third, note that Biden, McCain, the current US President Obama (Dem./Illinois) and US Secretary of State Hillary Clinton (Dem./New York) have very few outgoing arcs; e.g., Obama only directly influences Feingold (Dem./Wisconsin), a prominent senior member with strongly liberal stands. One may wonder why do such prominent senators seem to have so little direct influence on others? A possible explanation is that US President Bush was about to complete hist second term (the maximum allowed). Both parties had *very long* presidential primaries. All those senators contended for the presidential candidacy within their parties. Hence, one may posit that those senators were focusing on running their campaigns and that their influence in the *day-to-day* business of congress was channeled through other prominent senior members of their parties.

joint-actions are PSNE and which ones are not. If our only goal were to find a game $\mathcal{G}$ in which all the given observed data is a PSNE, then a "dummy" LIG with $\mathcal{G} = (\mathbf{W}, \mathbf{b}), \mathbf{W} = \mathbf{0}, \mathbf{b} = \mathbf{0}$ would be the optimal solution since $|\mathcal{NE}(\mathcal{G})| = 2^n$. In this section, we present a probabilistic formulation that allows finding games that maximize the *empirical proportion of equilibria* in the data while keeping the *true proportion of equilibria* low. Furthermore, we show that *trivial* games such as $\mathbf{W} = \mathbf{0}, \mathbf{b} = \mathbf{0}$, obtain the lowest log-likelihood.

**On the Identifiability of Games.** Several games with different coefficients can lead to the same PSNE set. It is easy to construct simple games with different structures or parameters, but the same set of PSNE. (Please see Appendix A for examples.) In this work, we choose to identify games by their PSNE sets.

**Definition 1.** *We say that two games $\mathcal{G}_1$ and $\mathcal{G}_2$ are* equivalent *if and only if their PSNE sets are identical, i.e.: $\mathcal{G}_1 \equiv_{\mathcal{NE}} \mathcal{G}_2 \Leftrightarrow \mathcal{NE}(\mathcal{G}_1) = \mathcal{NE}(\mathcal{G}_2)$.*

### 2.1 A Generative Model of Behavioral Data

We propose the following generative model for behavioral data based strictly in the context of "simultaneous"/one-shot play in non-cooperative game theory. Let $\mathcal{G}$ be a game. With some probability $0 < q < 1$, a joint-action $\mathbf{x}$ is chosen uniformly at random from $\mathcal{NE}(\mathcal{G})$; otherwise, $\mathbf{x}$ is chosen uniformly at random from its complement set $\{-1, +1\}^n - \mathcal{NE}(\mathcal{G})$. Hence, the generative model is a mixture model with mixture parameter $q$ corresponding to the probability that a stable outcome (i.e. a PSNE) of the game is observed. Formally, the probability mass function (PMF) over joint-behaviors $\{-1, +1\}^n$ parametrized by $(\mathcal{G}, q)$ is:

$$p_{(\mathcal{G},q)}(\mathbf{x}) = q\,\frac{1[\mathbf{x} \in \mathcal{NE}(\mathcal{G})]}{|\mathcal{NE}(\mathcal{G})|} + (1-q)\,\frac{1[\mathbf{x} \notin \mathcal{NE}(\mathcal{G})]}{2^n - |\mathcal{NE}(\mathcal{G})|} \tag{1}$$

where we can think of $q$ as the "signal" level, and thus $1 - q$ as the "noise" level in the data set.

(A more sophisticated noise process, and mixed-strategy Nash equilibria are discussed in Appendix G.)

**Remark 2.** *In order for eq.(1) to be a valid PMF for any $\mathcal{G}$, we need to enforce $|\mathcal{NE}(\mathcal{G})| = 0 \Rightarrow q = 0$ and $|\mathcal{NE}(\mathcal{G})| = 2^n \Rightarrow q = 1$. In both cases ($|\mathcal{NE}(\mathcal{G})| \in \{0, 2^n\}$) the PMF becomes a uniform distribution. On the other hand, if $0 < |\mathcal{NE}(\mathcal{G})| < 2^n$ then setting $q \in$*

$\{0, 1\}$ *leads to an invalid PMF.*

Let $\pi(\mathcal{G})$ be the *true proportion of equilibria* in the game $\mathcal{G}$ relative to all possible joint-actions, i.e.:

$$\pi(\mathcal{G}) \equiv |\mathcal{NE}(\mathcal{G})|/2^n \quad (2)$$

We say that a game $\mathcal{G}$ is *trivial* if and only if $|\mathcal{NE}(\mathcal{G})| \in \{0, 2^n\}$ (or equivalently $\pi(\mathcal{G}) \in \{0, 1\}$); and say $\mathcal{G}$ is *non-trivial* otherwise.

The following propositions establish that the condition $q > \pi(\mathcal{G})$ ensures that the probability of a PSNE is strictly greater than a non-PSNE, and it also guarantees identifiability among non-trivial games.

**Proposition 3.** *Given a non-trivial game $\mathcal{G}$, the mixture parameter $q > \pi(\mathcal{G})$ if and only if $p_{(\mathcal{G},q)}(\mathbf{x}_1) > p_{(\mathcal{G},q)}(\mathbf{x}_2)$ for any $\mathbf{x}_1 \in \mathcal{NE}(\mathcal{G})$ and $\mathbf{x}_2 \notin \mathcal{NE}(\mathcal{G})$.*

*Proof Sketch.* By eq.(1) and eq.(2). □

(Please, see Appendix B for detailed proofs.)

**Proposition 4.** *Let $\mathcal{G}_1$ and $\mathcal{G}_2$ be two non-trivial games. For some mixture parameter $q > \max(\pi(\mathcal{G}_1), \pi(\mathcal{G}_2))$, $\mathcal{G}_1$ and $\mathcal{G}_2$ are* equivalent *if and only if they induce the same PMF over the joint-action space $\{-1, +1\}^n$ of the players, i.e.: $\mathcal{G}_1 \equiv_{\mathcal{NE}} \mathcal{G}_2 \Leftrightarrow (\forall \mathbf{x})\ p_{(\mathcal{G}_1,q)}(\mathbf{x}) = p_{(\mathcal{G}_2,q)}(\mathbf{x})$.*

*Proof Sketch.* By Definition 1, eq.(1) and eq.(2). □

### 2.2 Learning Games via MLE

The *learning problem* consists on estimating the structure and parameters of a graphical game from data. We point out that our problem is unsupervised, i.e. we do not know a priori which joint-actions are PSNE and which ones are not.

First, we use a shorthand notation for the Kullback-Leibler (KL) divergence between two Bernoulli distributions parametrized by $0 \le p_1, p_2 \le 1$:

$$\begin{aligned} KL(p_1 \| p_2) &\equiv KL(\text{Bernoulli}(p_1) \| \text{Bernoulli}(p_2)) \\ &= p_1 \log \tfrac{p_1}{p_2} + (1 - p_1) \log \tfrac{1-p_1}{1-p_2} \end{aligned} \quad (3)$$

Next, we present the MLE problem for games.

**Lemma 5.** *Given a dataset $\mathcal{D} = \mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(m)}$, let $\widehat{\pi}(\mathcal{G})$ be the* empirical proportion of equilibria*, i.e. the proportion of samples in $\mathcal{D}$ that are PSNE of $\mathcal{G}$:*

$$\widehat{\pi}(\mathcal{G}) \equiv \tfrac{1}{m} \textstyle\sum_l 1[\mathbf{x}^{(l)} \in \mathcal{NE}(\mathcal{G})] \quad (4)$$

*the MLE problem for the probabilistic model in eq.(1) can be expressed as:*

$$\max_{(\mathcal{G},q)\in\Upsilon} \widehat{\mathcal{L}}(\mathcal{G}, q) = KL(\widehat{\pi}(\mathcal{G}) \| \pi(\mathcal{G})) - KL(\widehat{\pi}(\mathcal{G}) \| q) - n \log 2 \quad (5)$$

*where $\mathcal{H}$ is the class of games of interest, $\Upsilon = \{(\mathcal{G}, q) \mid \mathcal{G} \in \mathcal{H} \wedge 0 < \pi(\mathcal{G}) < q < 1\}$ is the hypothesis space of non-trivial identifiable games, $\pi(\mathcal{G})$ is defined as in eq.(2) and the optimal mixture parameter $\widehat{q} = \min(\widehat{\pi}(\mathcal{G}), 1 - \frac{1}{2m})$.*

*Proof Sketch.* The expression $\widehat{\mathcal{L}}(\mathcal{G}, q)$ follows from the definition of log-likelihood and algebraic manipulation for using eq.(3). By maximizing with respect to $q$, we get $KL(\widehat{\pi}(\mathcal{G}) \| \widehat{q}) = 0 \Leftrightarrow \widehat{q} = \widehat{\pi}(\mathcal{G})$. We define our hypothesis space $\Upsilon$ given the conditions in Remark 2 and Propositions 3 and 4. For $\widehat{\pi}(\mathcal{G}) = 1$, we "shrink" $\widehat{q}$ to $1 - \frac{1}{2m}$ in order to generate a valid PMF (Remark 2). □

**Remark 6.** *A trivial game (e.g. $\mathcal{G} = (\mathbf{W}, \mathbf{b}), \mathbf{W} = \mathbf{0}, \mathbf{b} = \mathbf{0}, \pi(\mathcal{G}) = 1$) induces a uniform PMF by Remark 2, and therefore its log-likelihood is $-n \log 2$. Note that the lowest log-likelihood for non-trivial identifiable games in eq.(5) is $-n \log 2$ by setting the optimal mixture parameter $\widehat{q} = \widehat{\pi}(\mathcal{G})$ and given that $KL(\widehat{\pi}(\mathcal{G}) \| \pi(\mathcal{G})) \ge 0$.*

## 3 Generalization Bound

In this section, we show a generalization bound for the MLE problem and an upper bound of the VC-dimension of LIGs. Our objective is to establish that with high probability, the maximum likelihood estimate is close to the optimal parameters, in terms of achievable expected log-likelihood.

Given the ground truth distribution $\mathcal{Q}$ of the data, let $\bar{\pi}(\mathcal{G})$ be the *expected proportion of equilibria*, i.e.:

$$\bar{\pi}(\mathcal{G}) = \mathbb{P}_{\mathcal{Q}}[\mathbf{x} \in \mathcal{NE}(\mathcal{G})]$$

and let $\bar{\mathcal{L}}(\mathcal{G}, q)$ be the *expected log-likelihood* of a generative model from game $\mathcal{G}$ and mixture parameter $q$, i.e.:

$$\bar{\mathcal{L}}(\mathcal{G}, q) = \mathbb{E}_{\mathcal{Q}}[\log p_{(\mathcal{G},q)}(\mathbf{x})]$$

Note that our hypothesis space $\Upsilon$ in eq.(5) includes a continuous parameter $q$ that could potentially have infinite VC-dimension. The following lemma will allow us later to prove that uniform convergence for the extreme values of $q$ implies uniform convergence for all $q$ in the domain.

**Lemma 7.** *Consider any game $\mathcal{G}$ and, for $0 < q'' < q' < q < 1$, let $\theta = (\mathcal{G}, q)$, $\theta' = (\mathcal{G}, q')$ and $\theta'' = (\mathcal{G}, q'')$. If, for any $\epsilon > 0$ we have $|\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta)| \le \epsilon/2$ and $|\widehat{\mathcal{L}}(\theta'') - \bar{\mathcal{L}}(\theta'')| \le \epsilon/2$, then $|\widehat{\mathcal{L}}(\theta') - \bar{\mathcal{L}}(\theta')| \le \epsilon/2$.*

*Proof Sketch.* By basic algebra, we have $\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta) = (\widehat{\pi}(\mathcal{G}) - \bar{\pi}(\mathcal{G})) \log \left( \frac{q}{1-q} \cdot \frac{1-\pi(\mathcal{G})}{\pi(\mathcal{G})} \right)$. Note that $\frac{q}{1-q}$ is strictly monotonically increasing for $0 \le q < 1$. □

The following theorem shows that the expected log-likelihood of the maximum likelihood estimate converges in probability to that of the optimal, as the data size $m$ increases.

**Theorem 8.** *Let $\widehat{\theta} = (\widehat{\mathcal{G}}, \widehat{q})$ be the maximum likelihood estimate in eq.(5) and $\bar{\theta} = (\bar{\mathcal{G}}, \bar{q})$ be the maximum expected likelihood estimate, i.e. $\widehat{\theta} = \arg\max_{\theta \in \Upsilon} \widehat{\mathcal{L}}(\theta)$ and $\bar{\theta} = \arg\max_{\theta \in \Upsilon} \bar{\mathcal{L}}(\theta)$, then with probability at least $1 - \delta$:*

$$\bar{\mathcal{L}}(\widehat{\theta}) \geq \bar{\mathcal{L}}(\bar{\theta}) - \left(\log \max(2m, \tfrac{1}{1-\bar{q}}) + n \log 2\right) \cdot \sqrt{\tfrac{2}{m}\left(\log d(\mathcal{H}) + \log \tfrac{4}{\delta}\right)}$$

*where $\mathcal{H}$ is the class of games of interest, $\Upsilon = \{(\mathcal{G}, q) \mid \mathcal{G} \in \mathcal{H} \wedge 0 < \pi(\mathcal{G}) < q < 1\}$ is the hypothesis space of non-trivial identifiable games and $d(\mathcal{H}) \equiv |\cup_{\mathcal{G} \in \mathcal{H}} \{\mathcal{NE}(\mathcal{G})\}|$ is the number of all possible games in $\mathcal{H}$ (identified by their PSNE sets).*

*Proof Sketch.* The log-likelihood is bounded and since $\mathbb{E}[\widehat{\mathcal{L}}(\theta)] = \bar{\mathcal{L}}(\theta)$, we use Hoeffding's inequality for each $\theta$. For applying the union bound for all $\theta$, note that there are $2d(\mathcal{H})$ possible parameters $\theta$, since by Lemma 7 we need to consider only the two extreme values of $q \in \{\pi(\mathcal{G}), \max(1 - \frac{1}{2m}, \bar{q})\}$. □

The following theorem establishes the complexity of the class of LIGs, which implies that $\log d(\mathcal{H})$ in Theorem 8 is polynomial in the number of players $n$.

**Theorem 9.** *Let $\mathcal{H}$ be the class of LIGs. Then $d(\mathcal{H}) \equiv |\cup_{\mathcal{G} \in \mathcal{H}} \{\mathcal{NE}(\mathcal{G})\}| \leq 2^{\frac{n^2(n+1)}{2}+1} \leq 2^{n^3}$ .*

*Proof Sketch.* For every LIG in $\mathcal{H}$, we define a neural network. Note that $\log d(\mathcal{H})$ is upper bounded by the VC-dimension of the class of those properly defined neural networks. Finally, we use the VC-dimension of neural networks [Sontag, 1998]. □

# 4 Algorithms

In this section, we approximate the MLE problem by maximizing the number of observed PSNE in the data, suitable for a class of games with small true proportion of equilibria. We then present our convex loss minimization approach. We also discuss baseline methods such as sigmoidal approximation and exhaustive search.

First, we discuss some negative results that justifies the use of simple approaches. Counting the number of PSNE is NP-hard for LIGs, and so is computing the log-likelihood function and therefore performing MLE. This is not a disadvantage relative to probabilistic graphical models, since computing the log-likelihood function is also NP-hard for Markov random fields, while learning is also NP-hard for Bayesian networks. General approximation techniques such as pseudo-likelihood estimation do not lead to tractable methods for learning LIGs. From an optimization perspective, the log-likelihood function is not continuous because of the number of PSNE. Furthermore, bounding the number of PSNE by known bounds for Ising models leads to trivial bounds. (Formal proofs and discussion are included in Appendix C.)

## 4.1 Exhaustive Search

As a first approach, consider solving the MLE problem in eq.(5) by an exact exhaustive search algorithm. Note that in a LIG each player separates hypercube vertices with a linear function, i.e. for $\mathbf{v} \equiv (\mathbf{w}_{i,-i}, b_i)$ and $\mathbf{y} \equiv (x_i \mathbf{x}_{-i}, -x_i) \in \{-1, +1\}^n$ we have $x_i({\mathbf{w}_{i,-i}}^{\mathrm{T}} \mathbf{x}_{-i} - b_i) = \mathbf{v}^{\mathrm{T}} \mathbf{y}$. Assume we assign a binary label to each vertex $\mathbf{y}$, then note that not all possible labelings are linearly separable. Labelings which are linearly separable are called *linear threshold functions (LTFs).* [Muroga, 1965] showed that the number of LTFs is at least $\alpha(n) \equiv 2^{0.33048 n^2}$. Therefore, an exhaustive search approach (for a single player) will take at least $\alpha(n)$ time. Unfortunately, enumerating all LTFs seems to be far from a trivial problem. By using results in Muroga [1971], a weight vector $\mathbf{v}$ with integer entries such that $(\forall i)\ |v_i| \leq \beta(n) \equiv (n+1)^{(n+1)/2}/2^n$ is sufficient to realize all possible LTFs. Therefore enumerating LIGs takes at most $(2\beta(n)+1)^{n^2} \approx (\frac{\sqrt{n+1}}{2})^{n^3}$ steps, and we propose the use of this method for $n \leq 4$. For $n = 4$ we found that the number of LIGs is 23,706.

## 4.2 From MLE to MEPE

We approximately perform MLE for LIGs, by solving a *maximum empirical proportion of equilibria (MEPE)* problem, i.e. by maximizing the PSNE in the observed data. This strategy allows us to avoid computing the (NP-hard) true proportion of equilibria. For a class of games with small true proportion of equilibria with high probability (such as LIGs as shown in Section 5), we use a lower bound of the log-likelihood with high probability. We also show that under very mild conditions, the parameters $(\mathcal{G}, q)$ belong to the hypothesis space of the original problem with high probability.

First, we derive bounds on the log-likelihood function.

**Lemma 10.** *Given a non-trivial game $\mathcal{G}$ with $0 < \pi(\mathcal{G}) < \widehat{\pi}(\mathcal{G})$, the KL divergence in the log-likelihood function in eq.(5) is bounded as follows:*

$$-\widehat{\pi}(\mathcal{G}) \log \pi(\mathcal{G}) - \log 2 < KL(\widehat{\pi}(\mathcal{G}) \| \pi(\mathcal{G})) < -\widehat{\pi}(\mathcal{G}) \log \pi(\mathcal{G})$$

*Proof Sketch.* From basic calculus arguments. □

The bounds are very informative when $\pi(\mathcal{G}) \to 0$ (or in our setting when $n \to +\infty$), since $\log 2$ becomes small when compared to $-\log \pi(\mathcal{G})$ (See Appendix B.7.1 for an illustration).

Next, we derive the MEPE problem from MLE.

**Theorem 11.** *Assume that with probability at least* $1-\delta$ *we have* $\pi(\mathcal{G}) \le \frac{\kappa^n}{\delta}$ *for* $0 < \kappa < 1$. *Maximizing a lower bound (with high probability) of the log-likelihood in eq.*(5) *is equivalent to maximizing the empirical proportion of equilibria:*

$$\max_{\mathcal{G} \in \mathcal{H}} \widehat{\pi}(\mathcal{G}) \tag{6}$$

*furthermore, for all games* $\mathcal{G}$ *such that* $\widehat{\pi}(\mathcal{G}) \ge \gamma$ *for some* $0 < \gamma < 1/2$, *for sufficiently large* $n > \log_\kappa(\delta\gamma)$ *and optimal mixture parameter* $\widehat{q} = \min(\widehat{\pi}(\mathcal{G}), 1 - \frac{1}{2m})$, *we have* $(\mathcal{G}, \widehat{q}) \in \Upsilon$, *where* $\Upsilon = \{(\mathcal{G}, q) \mid \mathcal{G} \in \mathcal{H} \wedge 0 < \pi(\mathcal{G}) < q < 1\}$ *is the hypothesis space of non-trivial identifiable games.*

*Proof Sketch.* For bounding the log-likelihood, we apply the lower bound in Lemma 10 in eq.(5). For proving $(\mathcal{G}, \widehat{q}) \in \Upsilon \Leftrightarrow 0 < \pi(\mathcal{G}) < \widehat{q} < 1$, we use $\pi(\mathcal{G}) \le \frac{\kappa^n}{\delta}$ and $\gamma \le \widehat{\pi}(\mathcal{G})$. □

### 4.3 Sigmoidal Approximation

A very simple optimization approach can be devised by using a sigmoid in order to approximate the 0/1 function $1[z \ge 0]$ in the MLE problem of eq.(5) and the MEPE problem of eq.(6). We use the following sigmoidal approximation:

$$1[z \ge 0] \approx H_{\alpha,\beta}(z) \equiv \tfrac{1}{2}(1 + \tanh(\tfrac{z}{\beta} - \operatorname{atanh}(1 - 2\alpha^{1/n}))) \tag{7}$$

The additional term $\alpha$ ensures that for $\mathcal{G} = (\mathbf{W}, \mathbf{b}), \mathbf{W} = \mathbf{0}, \mathbf{b} = \mathbf{0}$ we get $1[\mathbf{x} \in \mathcal{NE}(\mathcal{G})] \approx H_{\alpha,\beta}(0)^n = \alpha$. We perform gradient ascent on these objective functions that have many local maxima. When maximizing the "sigmoidal" likelihood, each step of the gradient ascent is NP-hard due to the "sigmoidal" true proportion of equilibria. Therefore, we propose the use of the sigmoidal MLE for $n \le 15$. In our implementation, we add an $\ell_1$-norm regularizer $-\rho\|\mathbf{W}\|_1$ where $\rho > 0$.

### 4.4 Convex Loss Minimization (CLM)

From an optimization perspective, it is more convenient to minimize a convex objective instead of a sigmoidal approximation with many local minima. Note that maximizing the empirical proportion of equilibria in eq.(6) is equivalent to minimizing the empirical proportion of non-equilibria, i.e. $\min_{\mathcal{G} \in \mathcal{H}} (1 - \widehat{\pi}(\mathcal{G}))$. Furthermore, $1 - \widehat{\pi}(\mathcal{G}) = \frac{1}{m}\sum_l 1[\mathbf{x}^{(l)} \notin \mathcal{NE}(\mathcal{G})]$. Denote by $\ell$ the 0/1 loss, i.e. $\ell(z) = 1[z < 0]$. For LIGs, solving the MEPE problem in eq.(6) is equivalent to solving the loss minimization problem:

$$\min_{\mathbf{W},\mathbf{b}} \frac{1}{m} \sum_l \max_i \ell(x_i^{(l)}(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i}^{(l)} - b_i)) \tag{8}$$

We can further relax this problem by introducing convex upper bounds of the 0/1 loss. As we will show in Section 5, the use of convex losses also avoids the trivial solution of eq.(8), i.e. $\mathbf{W} = \mathbf{0}, \mathbf{b} = \mathbf{0}$ (which obtains the lowest log-likelihood as discussed in Remark 6). In what follows, we develop four efficient methods for solving eq.(8). In our implementation, we add an $\ell_1$-norm regularizer $\rho\|\mathbf{W}\|_1$ where $\rho > 0$.

**Independent Support Vector Machines and Logistic Regression.** We can relax the loss minimization problem in eq.(8) by using the loose bound $\max_i \ell(z_i) \le \sum_i \ell(z_i)$. This relaxation simplifies the original problem into several independent problems. For each player $i$, we train the weights $(\mathbf{w}_{i,-i}, b_i)$ in order to predict independent (disjoint) actions. This leads to *1-norm SVMs* of Bradley and Mangasarian [1998], Zhu et al. [2003] and $\ell_1$-regularized logistic regression. We solve the latter with the *$\ell_1$-projection method* of Schmidt et al. [2007]. While the training is independent, our goal is not the prediction for independent players but the characterization of joint-actions. The use of these well known techniques in our context is novel, since we interpret the output of SVMs and logistic regression as the parameters of a LIG. Thus, we use the parameters to measure empirical and true proportion of equilibria, KL divergence and log-likelihood in our probabilistic model.

**Simultaneous Support Vector Machines.** While converting the loss minimization problem in eq.(8) by using loose bounds allow to obtain several independent problems with small number of variables, a second reasonable strategy would be to use tighter bounds at the expense of obtaining a single optimization problem with a higher number of variables.

For the hinge loss $\ell(z) = \max(0, 1-z)$, we have $\max_i \ell(z_i) = \max(0, 1-z_1, \dots, 1-z_n)$ and the loss minimization problem in eq.(8) becomes the following primal linear program:

$$\min_{\mathbf{W},\mathbf{b},\boldsymbol{\xi}} \frac{1}{m} \sum_l \xi_l + \rho\|\mathbf{W}\|_1$$
$$\text{s.t.} \ (\forall l, i)\ x_i^{(l)}(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i}^{(l)} - b_i) \ge 1 - \xi_l \ \ , \ \ (\forall l)\ \xi_l \ge 0 \tag{9}$$

where $\rho > 0$. This problem is a generalization of *1-norm SVMs* of Bradley and Mangasarian [1998], Zhu et al. [2003]. (The dual problem is in Appendix D.)

**Simultaneous Logistic Regression.** For the logistic loss $\ell(z) = \log(1 + e^{-z})$, we could use the nonsmooth loss $\max_i \ell(z_i)$ directly. Instead, we chose a smooth upper bound, i.e. $\log(1+\sum_i e^{-z_i})$ (Discussion is included in Appendix E.) The loss minimization problem in eq.(8) becomes:

$$\min_{\mathbf{W},\mathbf{b}} \frac{1}{m}\sum_l \log(1+\textstyle\sum_i e^{-x_i^{(l)}(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i}^{(l)}-b_i)}) + \rho\|\mathbf{W}\|_1 \tag{10}$$

where $\rho > 0$. We use the *$\ell_1$-projection method* of Schmidt et al. [2007] for optimizing eq.(10).

## 5 True Proportion of Equilibria

In this section, we define *absolute indifference* of players and show that convex loss minimization produces games in which all players are non-absolutely-indifferent. We then provide a bound of the true proportion of equilibria with high probability. Our bound only assumes independence of weight vectors among players. Our bound is distribution-free, i.e. we do not assume a specific distribution for the weight vector of each player. Furthermore, we do not assume any connectivity properties of the underlying graph.

Parallel to our work, Daskalakis et al. [2011] analyzed a different setting: random games which structure is drawn from the Erdős-Rényi model (i.e. each edge is present independently with the same probability $p$) and utility functions which are random tables. The analysis in Daskalakis et al. [2011], while more general than ours (which only focus on LIGs), it is at the same time more restricted since it assumes either the Erdős-Rényi model for random structures or connectivity properties for deterministic structures.

### 5.1 CLM Produces Games with Non-Absolutely-Indifferent Players

First, we define the notion of *absolute indifference* of players. Our goal is to show that our proposed convex loss algorithms produce LIGs in which all players are non-absolutely-indifferent and therefore every player defines constraints to the true proportion of equilibria.

**Definition 12.** *Given a LIG $\mathcal{G} = (\mathbf{W},\mathbf{b})$, we say a player $i$ is* absolutely indifferent *if and only if $(\mathbf{w}_{i,-i}, b_i) = \mathbf{0}$, and* non-absolutely-indifferent *if and only if $(\mathbf{w}_{i,-i}, b_i) \neq \mathbf{0}$.*

Next, we show that independent and simultaneous SVM and logistic regression produce games in which all players are non-absolutely-indifferent except for some "degenerate" cases. The following lemma applies to independent SVMs for $c^{(l)} = 0$ and simultaneous SVMs for $c^{(l)} = \max(0, \max_{j\neq i}(1 - x_j^{(l)}(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i}^{(l)} - b_i)))$.

**Lemma 13.** *Given $(\forall l)\ c^{(l)} \geq 0$, the minimization of the hinge training loss $\widehat{\ell}(\mathbf{w}_{i,-i}, b_i) = \frac{1}{m}\sum_l \max(c^{(l)}, 1 - x_i^{(l)}(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i}^{(l)} - b_i))$ guarantees non-absolutely-indifference of player $i$ except for some "degenerate" cases, i.e. the optimal solution $(\mathbf{w}^*_{i,-i}, b^*_i) = \mathbf{0}$ if and only if $(\forall j \neq i)\ \sum_l 1[x_i^{(l)}x_j^{(l)} = 1]u^{(l)} = \sum_l 1[x_i^{(l)}x_j^{(l)} = -1]u^{(l)}$ and $\sum_l 1[x_i^{(l)} = 1]u^{(l)} = \sum_l 1[x_i^{(l)} = -1]u^{(l)}$ where $u^{(l)}$ is defined as $c^{(l)} > 1 \Leftrightarrow u^{(l)} = 0$, $c^{(l)} < 1 \Leftrightarrow u^{(l)} = 1$ and $c^{(l)} = 1 \Leftrightarrow u^{(l)} \in [0;1]$.*

*Proof Sketch.* By optimality arguments, the subgradient of $\widehat{\ell}$ "vanishes" at $(\mathbf{w}^*_{i,-i}, b^*_i) = \mathbf{0}$. □

**Remark 14.** *For independent SVMs, the "degenerate" cases in Lemma 13 simplify to $(\forall j \neq i)\ \sum_l 1[x_i^{(l)}x_j^{(l)} = 1] = \frac{m}{2}$ and $\sum_l 1[x_i^{(l)} = 1] = \frac{m}{2}$.*

The following lemma applies to independent logistic regression for $c^{(l)} = 0$ and simultaneous logistic regression for $c^{(l)} = \sum_{j\neq i} e^{-x_j^{(l)}(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i}^{(l)} - b_i)}$.

**Lemma 15.** *Given $(\forall l)\ c^{(l)} \geq 0$, the minimization of the logistic training loss $\widehat{\ell}(\mathbf{w}_{i,-i}, b_i) = \frac{1}{m}\sum_l \log(c^{(l)} + 1 + e^{-x_i^{(l)}(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i}^{(l)} - b_i)})$ guarantees non-absolutely-indifference of player $i$ except for some "degenerate" cases, i.e. the optimal solution $(\mathbf{w}^*_{i,-i}, b^*_i) = \mathbf{0}$ if and only if $(\forall j \neq i)\ \sum_l \frac{1[x_i^{(l)}x_j^{(l)}=1]}{c^{(l)}+2} = \sum_l \frac{1[x_i^{(l)}x_j^{(l)}=-1]}{c^{(l)}+2}$ and $\sum_l \frac{1[x_i^{(l)}=1]}{c^{(l)}+2} = \sum_l \frac{1[x_i^{(l)}=-1]}{c^{(l)}+2}$.*

*Proof Sketch.* By optimality arguments, the gradient of $\widehat{\ell}$ vanishes at $(\mathbf{w}^*_{i,-i}, b^*_i) = \mathbf{0}$. □

**Remark 16.** *For independent logistic regression, the "degenerate" cases in Lemma 15 simplify to $(\forall j \neq i)\ \sum_l 1[x_i^{(l)}x_j^{(l)} = 1] = \frac{m}{2}$ and $\sum_l 1[x_i^{(l)} = 1] = \frac{m}{2}$.*

Based on these results, after termination of our proposed algorithms, we fix cases in which the optimal solution $(\mathbf{w}^*_{i,-i}, b^*_i) = \mathbf{0}$ by setting $b^*_i = 1$ if the action of player $i$ was mostly $-1$ or $b^*_i = -1$ otherwise. Note that our proofs still hold if we include the $\ell_1$-regularization term since the subdifferential of $\rho\|\mathbf{w}_{i,-i}\|_1$ vanishes at $\mathbf{w}_{i,-i} = 0$.

### 5.2 Bounding the True Proportion of PSNE

In what follows, we show that for a game with a single *non-absolutely-indifferent* player, the true proportion of equilibria is bounded by 3/4.

**Lemma 17.** *Given a LIG $\mathcal{G} = (\mathbf{W},\mathbf{b})$ with non-absolutely-indifferent player $i$ and absolutely-indifferent players $\forall j \neq i$, the following statements hold:*

i. $\mathbf{x} \in \mathcal{NE}(\mathcal{G}) \Leftrightarrow x_i(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i} - b_i) \geq 0$
ii. $|\mathcal{NE}(\mathcal{G})| = 2^{n-1} + \sum_{\mathbf{x}_{-i}} 1[\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i} - b_i = 0]$
iii. $\frac{1}{2} \leq \pi(\mathcal{G}) \leq \frac{3}{4}$

*Proof Sketch.* Claims i and ii follow from basic algebra. For proving Claim iii, we use results in Aichholzer and Aurenhammer [1996] regarding the number of vertices of a hypercube that are covered by a hyperplane. □

Next, we present our bound for the true proportion of equilibria of games in which all players are non-absolutely-indifferent.

**Theorem 18.** *If all players are non-absolutely-indifferent and if the rows of a LIG* $\mathcal{G} = (\mathbf{W}, \mathbf{b})$ *are independent (but not necessarily identically distributed) random vectors, i.e. for every player* $i$, $(\mathbf{w}_{i,-i}, b_i)$ *is independently drawn from an arbitrary distribution* $\mathcal{P}_i$*, then the expected true proportion of equilibria is bounded as follows:*

$$(1/2)^n \leq \mathbb{E}_{\mathcal{P}_1,\dots,\mathcal{P}_n}[\pi(\mathcal{G})] \leq (3/4)^n$$

*furthermore, the following high probability statement holds:*

$$\mathbb{P}_{\mathcal{P}_1,\dots,\mathcal{P}_n}[\pi(\mathcal{G}) \leq \tfrac{(3/4)^n}{\delta}] \geq 1 - \delta$$

*Proof Sketch.* We use Lemma 17 and the fact that functions of independent variables are independent. The high probability statement follows from Markov's inequality. □

## 6 Experimental Results

For learning LIGs we used our convex loss methods: independent and simultaneous SVM and logistic regression. We also used the (super-exponential) exhaustive search method for $n \leq 4$. As a baseline, we used the sigmoidal (NP-hard) MLE for $n \leq 15$ and the sigmoidal MEPE. We found experimentally that the parameters $\alpha = 0.1$ and $\beta = 0.001$ in the sigmoidal function achieved the best results.

For comparison, we learn Ising models. For $n \leq 15$ players, we perform exact $\ell_1$-regularized MLE by using FOBOS [Duchi and Singer, 2009] and exact (NP-hard) gradients. For $n > 15$ players, we use the method of [Höfling and Tibshirani, 2009] which uses a sequence of first-order approximations of the exact log-likelihood. We also used a two-step algorithm, by first learning the structure by $\ell_1$-regularized logistic regression [Wainwright et al., 2006] and then using FOBOS with belief propagation for gradient approximation. We did not find a statistically significant difference between both algorithms and reported the latter.

Our experimental setup is as follows: after learning models for different regularization levels $\rho$ in a training set, we select the value of $\rho$ that maximizes the log-likelihood in a validation set, and report statistics in a test set. For synthetic experiments, we report the KL divergence, average precision (one minus the fraction of falsely included PSNE), average recall (one minus the fraction of falsely excluded PSNE). For real-world experiments, we report the log-likelihood. For both, we report the number of PSNE and the empirical proportion of equilibria.

**Synthetic Experiments.** We use a small synthetic model in order to compare with the (super-exponential) exhaustive search method. The ground truth model $\mathcal{G}_g = (\mathbf{W}_g, \mathbf{b}_g)$ has $n = 4$ players and 4 PSNE, $\mathbf{W}_g$ was set according to Figure 2 (edge weights were set to $+1$) and $\mathbf{b}_g = \mathbf{0}$. The mixture parameter of the ground truth $q_g$ was set to 0.5,0.7,0.9. For each of 50 repetitions, we generated a training, a validation and a test set of 50 samples each. Figure 2 shows that our convex loss methods and sigmoidal MLE outperform (lower KL) exhaustive search, sigmoidal MEPE and Ising models. Exhaustive search (exact MLE) suffers from over-fitting since it does not obtain the lowest KL. From all convex loss methods, simultaneous logistic regression achieves the lowest KL. For all methods, the recovery of PSNE is perfect for $q_g = 0.9$ (equilibrium precision and recall equal to 1). The empirical proportion of equilibria resembles the ground truth mixture parameter $q_g$.

Next, we show that the performance of convex loss minimization improves as the number of samples increases. We use larger random models with varying number of samples (10,30,100,300). The ground truth model $\mathcal{G}_g = (\mathbf{W}_g, \mathbf{b}_g)$ contains $n = 20$ players. For each of 20 repetitions, we generate edges in $\mathbf{W}_g$ with density 0.5. Edge weights were set to $+1$ with probability $P(+1)$ and to $-1$ with probability $1 - P(+1)$. We set the bias $\mathbf{b}_g = \mathbf{0}$ and the mixture parameter of the ground truth $q_g = 0.7$. We then generated a training and a validation set with the same number of samples. Figure 3(a) shows that our convex loss methods outperform (lower KL) sigmoidal MEPE and Ising models. From all convex loss methods, simultaneous logistic regression achieves the lowest KL.

By using the above synthetic model, we evaluate the effect of our approximation methods. First, we test the impact of removing the true proportion of equilibria from our objective function. Second, we test the impact of using convex losses instead of a sigmoidal approximation of the 0/1 loss. We use a varying number of players ($n = 4, 6, 8, 10, 12$) and 50 samples. Figure 3(b) shows that in general, convex loss methods outperform (lower KL) sigmoidal MEPE, and MEPE outperforms sigmoidal MLE.

**Congressional Voting.** We used the U.S. congressional voting records in order to measure the generalization performance of convex loss minimization in a

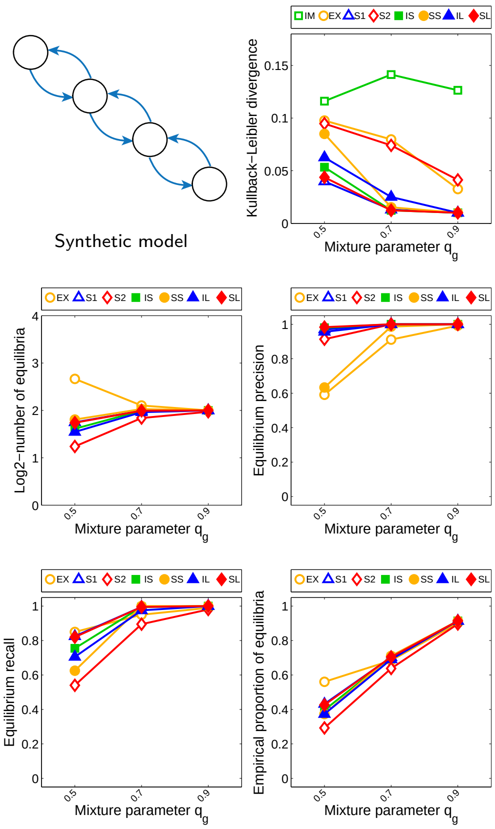


Figure 2: Closeness of recovered models to the ground truth for different $q_g$. Convex loss methods (IS,SS: independent and simultaneous SVM; IL,SL: logistic regression) and sigmoidal MLE (S1) have lower KL than exhaustive search (EX), sigmoidal MEPE (S2) and Ising models (IM). The recovery of PSNE is perfect for $q_g = 0.9$ and the empirical proportion of equilibria resembles $q_g$.

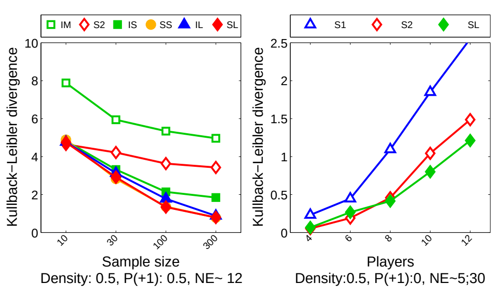


(a) (b)

Figure 3: KL divergence between recovered models and ground truth for increasing samples (a) and players (b), for 0.5 density. Each chart shows the probability $P(+1)$ that an edge has weight $+1$, and average number of equilibria (NE). In (a) convex loss methods (IS,SS: independent and simultaneous SVM; IL,SL: logistic regression) have lower KL than sigmoidal MEPE (S2) and Ising models (IM). In (b) simultaneous logistic regression (SL) has lower KL than sigmoidal MEPE (S2), and MEPE has lower KL than sigmoidal MLE (S1). (Omitted convex losses behave as SL).

real-world dataset. The dataset is publicly available at `http://www.senate.gov/` and covers from the 101th congress to the 111th congress (Jan 1989 to Dec 2010). The number of votes casted for each session were average: 337, minimum: 215, maximum: 613. Abstentions were replaced with negative votes. Since reporting the log-likelihood requires computing the (NP-hard) number of PSNE, we selected only 20 senators by stratified random sampling. We randomly split the data into three parts. We performed six repetitions by making each third of the data take turns as training, validation and testing sets. Figure 4 shows that our convex loss methods outperform (higher log-likelihood) sigmoidal MEPE and Ising models. From all convex loss methods, simultaneous logistic regression achieves the best KL. For all methods, the number of PSNE is low.

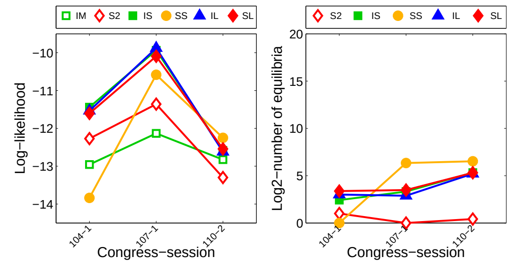


Figure 4: Statistics for games learnt from 20 senators, 1st session of the 104th and 107th congress and 2nd session of the 110th congress. The log-likelihood of convex loss methods (IS,SS: independent and simultaneous SVM; IL,SL: logistic regression) is higher than sigmoidal MEPE (S2) and Ising models (IM). The number of PSNE is low.

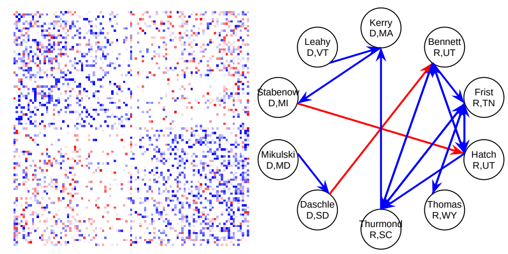


(a) (b)

Figure 5: Matrix of influence weights for games learnt from 100 senators, 1st session of the 107th congress, by using simultaneous logistic regression. A row represents how every other senator influences the senator in the row. Positive/negative influences in blue/red. Democrats in the top/left corner, Republicans in the bottom/right. Partial view of the graph (b).

We apply convex loss minimization to larger problems, by learning structures of games from all 100 senators. Figure 5 shows that simultaneous logistic regression elicits the bipartisan structure of the congress.

(Please, see Appendix F for additional results.)

**Concluding Remarks.** There are several ways of extending this research. More sophisticated noise processes as well as mixed-strategy Nash equilibria need to be considered and studied. Finally, topic-specific and time-varying versions of our model would elicit differences in preferences and trends.

(We include additional discussion in Appendix G.)

## References


O. Aichholzer and F. Aurenhammer. Classifying Hyperplanes in Hypercubes. *SIAM Journal on Discrete Mathematics*, 9(2):225–232, 1996.

R. Aumann. Subjectivity and Correlation in Randomized Strategies. *Journal of Mathematical Economics*, 1, 1974.

B. Blum, C.R. Shelton, and D. Koller. A Continuation Method for Nash Equilibria in Structured Games. *JAIR*, 25:457–502, 2006.

P. Bradley and O. Mangasarian. Feature Selection Via Concave Minimization and Support Vector Machines. *ICML*, 1998.

C. Daskalakis, A. Dimakisy, and E. Mossel. Connectivity and Equilibrium in Random Games. *Annals of Applied Probability*, 21(3):987–1016, 2011.

C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. *Commun. ACM*, 52(2):89–97, 2009.

J. Duchi and Y. Singer. Efficient Learning using Forward-Backward Splitting. *NIPS*, 2009.

Q. Duong, Y. Vorobeychik, S. Singh, and M. Wellman. Learning Graphical Game Models. *IJCAI*, 2009.

S. Ficici, D. Parkes, and A. Pfeffer. Learning and Solving Many-Player Games through a Cluster-Based Representation. *UAI*, 2008.

D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.

X. Gao and A. Pfeffer. Learning Game Representations from Data Using Rationality Constraints. *UAI*, 2010.

H. Höfling and R. Tibshirani. Estimation of Sparse Binary Pairwise Markov Networks using Pseudo-likelihoods. *JMLR*, 2009.

M. Irfan and L. Ortiz. A Game-Theoretic Approach to Influence in Networks. *AAAI*, 2011.

A. Jiang and K. Leyton-Brown. Polynomial-time Computation of Exact Correlated Equilibrium in Compact Games. *ACM Electronic Commerce Conference*, 2011.

S. Kakade, M. Kearns, J. Langford, and L. Ortiz. Correlated equilibria in graphical games. *ACM Electronic Commerce Conference*, 2003.

M. Kearns, M. Littman, and S. Singh. Graphical Models for Game Theory. *UAI*, 2001.

M. Kearns and J. Wortman. Learning from Collective Behavior. *COLT*, 2008.

S. Muroga. Lower bounds on the number of threshold functions and a maximum weight. *IEEE Transactions on Electronic Computers*, 14:136148, 1965.

S. Muroga. *Threshold Logic and Its Applications*. John Wiley & Sons, 1971.

J. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

L. Ortiz and M. Kearns. Nash Propagation for Loopy Graphical Games. *NIPS*, 2002.

C. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3):1–29, 2008.

M. Schmidt, G. Fung, and R. Rosales. Fast Optimization Methods for $\ell_1$ Regularization: A Comparative Study and Two New Approaches. *ECML*, 2007.

E. Sontag. VC Dimension of Neural Networks. In *Neural Networks and Machine Learning*, pages 69–95. Springer, 1998.

D. Vickrey and D. Koller. Multi-Agent Algorithms for Solving Graphical Games. *AAAI*, 2002.

Y. Vorobeychik, M. Wellman, and S. Singh. Learning Payoff Functions in Infinite Games. *IJCAI*, 2005.

M. Wainwright, P. Ravikumar, and J. Lafferty. High dimensional Graphical Model Selection Using $\ell_1$-Regularized Logistic Regression. *NIPS*, 2006.

K. Waugh, B. Ziebart, and J. Bagnell. Computational Rationalization: The Inverse Equilibrium Problem. *ICML*, 2011.

J. Wright and K. Leyton-Brown. Beyond Equilibrium: Predicting Human Behavior in Normal Form Games. *AAAI*, 2010.

J. Zhu, S. Rosset, T. Hastie, and R. Tibshirani. 1-norm Support Vector Machines. *NIPS*, 2003.

B. Ziebart, J. Bagnell, and A. Dey. Modeling Interaction via the Principle of Maximum Causal Entropy. *ICML*, 2010.